



(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
06.02.2002 Bulletin 2002/06

(51) Int Cl.7: G10L 17/00

(21) Application number: 01305725.2

(22) Date of filing: 02.07.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

- Thyes, Oliver
2551 Luxemburg (LU)
- Nguyen, Patrick
Santa Barbara, California 93103 (US)
- Junqua, Jean-Claude
Santa Barbara, California 93111 (US)
- Boman, Robert
Thousand Oaks, California 91360 (US)

(30) Priority: 05.07.2000 US 610495

(71) Applicant: Matsushita Electric Industrial Co., Ltd.
Kadoma-shi Osaka-fu, 571 (JP)

(72) Inventors:
• Kuhn, Roland
Santa Barbara, California 93111 (US)

(74) Representative:
Burrington, Alan Graham Headford et al
Atkinson Burrington 25-29 President Buildings
President Way
Sheffield S4 7UR (GB)

(54) Speaker verification and identification in their own spaces

(57) Client speaker locations in a speaker space are used to generate speech models for comparison with test speaker data or test speaker speech models. The speaker space can be constructed using training speakers that are entirely separate from the population of client speakers, or from client speakers, or from a mix of training and client speakers. Reestimation of the speaker space based on client environment information is also provided to improve the likelihood that the client data will fall within the speaker space. During enrollment of the clients into the speaker space, additional client speech can be obtained when predetermined conditions are met. The speaker distribution can also be used in the client enrollment step.




FIG. 3

Description**Cross-Reference to Related Application**

- 5 [0001] This is continuation-in-part of U.S. application S.N. 09/148,911 filed September 4, 1998, entitled "Speaker Verification and Speaker Identification Based on Eigenvoices."

Background and Summary of the Invention

- 10 [0002] The present invention relates generally to speech technology and, more particularly, to a system and method for performing speaker verification or speaker identification.

- [0003] The problem of authentication lies at the heart of nearly every transaction. Millions of people conduct confidential financial transactions over the telephone, such as accessing their bank accounts or using their credit cards. Authentication under current practice is far from foolproof. The parties exchange some form of presumably secret information, such as social security number, mother's maiden name or the like. Clearly, such information can be pirated, resulting in a false authentication.

- [0004] One aspect of the present invention addresses the foregoing problem by providing a system and method for performing speaker verification. Speaker verification involves determining whether a given voice belongs to a certain speaker (herein called the "client") or to an impostor (anyone other than the client).

- 20 [0005] Somewhat related to the problem of speaker verification is the problem of speaker identification. Speaker identification involves matching a given voice to one of a set of known voices. Like speaker verification, speaker identification has a number of attractive applications. For example, a speaker identification system may be used to classify voice mail by speaker for a set of speakers for which voice samples are available. Such capability would allow a computer-implemented telephony system to display on a computer screen the identity of callers who have left messages on the voice mail system.

- [0006] While the applications for speaker verification and speaker identification are virtually endless, the solution to performing these two tasks has heretofore proven elusive. Recognizing human speech and particularly discriminating the speaker from other speakers is a complex problem. Rarely does a person speak even a single word the same way twice due to how human speech is produced.

- 30 [0007] Human speech is the product of air under pressure from the lungs being forced through the vocal cords and modulated by the glottis to produce sound waves that then resonate in the oral and nasal cavities before being articulated by the tongue, jaw, teeth and lips. Many factors affect how these sound producing mechanisms inter-operate. The common cold, for example, greatly alters the resonance of the nasal cavity as well as the tonal quality of the vocal cords.

- 35 [0008] Given the complexity and variability with which the human produces speech, speaker verification and speaker identification are not readily performed by comparing new speech with a previously recorded speech sample. Employing a high similarity threshold, to exclude impostors, may exclude the authentic speaker when he or she has a head cold. On the other hand, employing a low similarity threshold can make the system prone to false verification.

- 40 [0009] The present invention uses a model-based analytical approach to speaker verification and speaker identification. Models are constructed and trained upon the speech of known client speakers (and possibly in the case of speaker verification also upon the speech of one or more impostors). These speaker models typically employ a multiplicity of parameters (such as Hidden Markov Model or GMM parameters). Rather than using these parameters directly, the parameters are concatenated to form supervectors. These supervectors, one supervector per speaker, represent the entire training data speaker population.

- 45 [0010] A linear transformation is performed on the supervectors resulting in a dimensionality reduction that yields a low-dimensional space that we call eigenspace. The basis vectors of this eigenspace we call "eigenvoice" vectors or "eigenvectors". If desired, the eigenspace can be further dimensionally reduced by discarding some of the eigenvector terms.

- 50 [0011] Next, each of the speakers comprising the training data is represented in eigenspace, either as a point in eigenspace or as a probability distribution in eigenspace. The former is somewhat less precise in that it treats the speech from each speaker as relatively unchanging. The latter reflects that the speech of each speaker will vary from utterance to utterance.

- [0012] Having represented the training data for each speaker in eigenspace, the system may then be used to perform speaker verification or speaker identification.

- 55 [0013] New speech data is obtained and used to construct a supervector that is then dimensionally reduced and represented in the eigenspace. Assessing the proximity of the new speech data to prior data in eigenspace, speaker verification or speaker identification is performed. The new speech from the speaker is verified if its corresponding point or distribution within eigenspace is within a threshold proximity to the training data for that client speaker. The

system may reject the new speech as authentic if it falls closer to an impostor's speech when placed in eigenspace.

[0014] Speaker identification is performed in a similar fashion. The new speech data is placed in eigenspace and identified with that training speaker whose eigenvector point for distribution is closest.

[0015] Assessing proximity between the new speech data and the training data in eigenspace has a number of advantages. First, the eigenspace represents in a concise, low-dimensional way, each entire speaker, not merely a selected few features of each speaker. Proximity computations performed in eigenspace can be made quite rapidly as there are typically considerably fewer dimensions to contend with in eigenspace than there are in the original speaker model space or feature vector space. Also, the system does not require that the new speech data include each and every example or utterance that was used to construct the original training data. Through techniques described herein, it is possible to perform dimensionality reduction on a supervector for which some of its components are missing. The result point for distribution in eigenspace nevertheless will represent the speaker remarkably well.

[0016] For a more complete understanding of the invention, its objects and advantages, refer to the following specification and to the accompanying drawings.

Brief Description of the Drawings

[0017]

Figure 1 illustrates an exemplary Hidden Markov Model (HMM), useful in understanding the invention;

Figure 2 is a flow diagram showing how the eigenspace may be constructed to implement a speaker identification system, where known client speakers are represented as points in eigenspace;

Figure 3 is a flow diagram illustrating how the eigenspace may be constructed to implement a speaker verification system, where the client speaker and potential impostors are represented as distributions in eigenspace;

Figure 4 is a flow diagram illustrating the process by which either speaker identification or speaker verification may be performed using the eigenspace developed during training;

Figure 5 is an illustration of how the maximum likelihood technique is performed;

Figure 6 is a data structure diagram illustrating how the observation data from a speaker may be placed into eigenspace based on the maximum likelihood operation;

Figure 7 illustrates an exemplary Gaussian Mixture Model (GMM), useful in understanding the invention;

Figure 8 is a flow diagram showing how points from the eigenspace may be used to generate probabilistic speech models;

Figure 9 is a flow diagram illustrating the process by which either speaker identification or speaker verification may be performed using speech models generated from the speaker space;

Figure 10 is a flow diagram showing one approach to enrolling speakers to solicit additional speech based on an average distance between spoken words; and

Figure 11 is a flow diagram showing another approach to enrolling speakers to solicit additional speech based on population density within the speaker space.

Description of the Preferred Embodiments

[0018] The eigenvoice techniques employed by the present invention will work with many different speech models. We illustrate the preferred embodiment in connection with a Hidden Markov Model recognizer because of its popularity in speech recognition technology today. However, it should be understood that the invention can be practiced using other types of model-based recognizers, such as phoneme similarity recognizers, for example.

[0019] To better understand the speaker identification and verification techniques of the invention, a basic understanding of speech recognition systems will be helpful. Inasmuch as most present day speech recognizers employ Hidden Markov Models (HMMs) to represent speech, the HMM technology will be described here to familiarize the reader.

[0020] The Hidden Markov Model is a modeling approach involving state diagrams. Any speech unit (such as a phrase, word, subword, phoneme or the like) can be modeled, with all knowledge sources included in that model. The HMM represents an unknown process that produces a sequence of observable outputs at discrete intervals, the outputs being members of some finite alphabet (corresponding to the predefined set of speech units). These models are called "hidden" because the state sequence that produced the observable output is not known,

[0021] As illustrated in Figure 1, an HMM 10 is illustrated by a set of states (S1, S2 ... S5), vectors that define transitions between certain pairs of states, illustrated as arrows in Figure 1, and a collection of probability data. Specifically, the Hidden Markov Model includes a set of transition probabilities 12 associated with the transition vectors and a set of output probabilities 14 associated with the observed output at each state. The model is clocked from one state to another at regularly spaced, discrete intervals. At clock-time, the model may change from its current state to any state

for which a transition vector exists. As illustrated, a transition can be from a given state back to itself.

[0022] The transition probabilities represent the likelihood that a transition from one state to another will occur when the model is clocked. Thus, as illustrated in Figure 1, each transition has associated with it a probability value (between 0 and 1). The sum of all probabilities leaving any state equals 1. For illustration purposes, a set of exemplary transition probability values has been given in transition probability Table 12. It will be understood that in a working embodiment these values would be generated by the training data, with the constraint that the sum of all probabilities leaving any state equals 1.

[0023] Every time a transition is taken, the model can be thought of as emitting or outputting one member of its alphabet. In the embodiment illustrated in Figure 1, a phoneme-based speech unit has been assumed. Thus the symbols identified in output probability Table 14 correspond to some of the phonemes found in standard English. Which member of the alphabet gets emitted upon each transition depends on the output probability value or function learned during training. The outputs emitted thus represent a sequence of observations (based on the training data) and each member of the alphabet has a probability of being emitted.

[0024] In modeling speech, it is common practice to treat the output as a sequence of continuous vectors as opposed to a sequence of discrete alphabet symbols. This requires the output probabilities to be expressed as continuous probability functions, as opposed to single numeric values. Thus HMMs are often based on probability functions comprising one or more Gaussian distributions. When a plurality of Gaussian functions are used they are typically additively mixed together to define a complex probability distribution, as illustrated at 16.

[0025] Whether represented as a single Gaussian function or a mixture of Gaussian functions, the probability distributions can be described by a plurality of parameters. Like the transition probability values (Table 12) these output probability parameters may comprise floating point numbers. Parameters Table 18 identifies the parameters typically used to represent probability density functions (pdf) based on observed data from the training speakers. As illustrated by the equation in Figure 1 at Gaussian function 16, the probability density function for an observation vector O to be modeled is the iterative sum of the mixture coefficient for each mixture component multiplied by the Gaussian density n , where the Gaussian density has a mean vector u_j and covariance matrix U_j computed from the cepstral or filter bank coefficient speech parameters.

[0026] The implementation details of a Hidden Markov Model recognizer may vary widely from one application to another. The HMM example shown in Figure 1 is intended merely to illustrate how Hidden Markov Models are constructed, and is not intended as a limitation upon the scope of the present invention. In this regard, there are many variations on the Hidden Markov Modeling concept. As will be more fully understood from the description below, the eigenvoice adaptation technique of the invention can be readily adapted to work with each of the different Hidden Markov Model variations, as well as with other parameter-based speech modeling systems.

[0027] Figures 2 and 3 illustrate, respectively, how speaker identification and speaker verification may be performed using the techniques of the invention. As a first step in performing either speaker identification or speaker verification, an eigenspace is constructed. The specific eigenspace constructed depends upon the application. In the case of speaker identification, illustrated in Figure 2, a set of known client speakers 20 is used to supply training data 22 upon which the eigenspace is created. Alternatively, for speaker verification, shown in Figure 3, the training data 22 are supplied from the client speaker or speakers 21a for which verification will be desired and also from one or more potential impostors 21b. Aside from this difference in training data source, the procedure for generating the eigenspace is essentially the same for both speaker identification and speaker verification applications. Accordingly, like reference numerals have been applied to Figures 2 and 3.

[0028] Referring to Figures 2 and 3, the eigenspace is constructed by developing and training speaker models for each of the speakers represented in the training data 22. This step is illustrated at 24 and generates a set of models 26 for each speaker. Although Hidden Markov Models have been illustrated here, the invention is not restricted to Hidden Markov Models. Rather, any speech model having parameters suitable for concatenation may be used. Preferably, the models 26 are trained with sufficient training data so that all sound units defined by the model are trained by at least one instance of actual speech for each speaker. Although not illustrated explicitly in Figures 2 and 3, the model training step 24 can include appropriate auxiliary speaker adaptation processing to refine the models. Examples of such auxiliary processing include Maximum A Posteriori estimation (MAP) or other transformation-based approaches such as Maximum Likelihood Linear Regression (MLLR). The objective in creating the speaker models 26 is to accurately represent the training data corpus, as this corpus is used to define the metes and bounds of the eigenspace into which each training speaker is placed and with respect to which each new speech utterance is tested.

[0029] After constructing the models 26, the models for each speaker are used to construct a supervector at step 28. The supervector, illustrated at 30, may be formed by concatenating the parameters of the model for each speaker. Where Hidden Markov Models are used, the supervector for each speaker may comprise an ordered list of parameters (typically floating point numbers) corresponding to at least a portion of the parameters of the Hidden Markov Models for that speaker. Parameters corresponding to each sound unit are included in the supervector for a given speaker. The parameters may be organized in any convenient order. The order is not critical; however, once an order is adopted

it must be followed for all training speakers.

[0030] The choice of model parameters to use in constructing the supervector will depend on the available processing power of the computer system. When using Hidden Markov Model parameters, we have achieved good results by constructing supervectors from the Gaussian means. If greater processing power is available, the supervectors may also include other parameters, such as the transition probabilities (Table 12, Fig. 1) or the Covariance Matrix parameters (parameters 18, Fig. 1). If the Hidden Markov Models generate discrete outputs (as opposed to probability densities), then these output values may be used to comprise the supervector.

[0031] After constructing the supervectors a dimensionality reduction operation is performed at step 32. Dimensionality reduction can be effected through any linear transformation that reduces the original high-dimensional supervectors into basis vectors. A non-exhaustive list of examples includes:

Principal Component Analysis (PCA), Independent Component Analysis (ICA),
Linear Discriminant Analysis (LDA)), Factor Analysis (FA), and Singular Value Decomposition (SVD).

[0032] More specifically, the class of dimensionality reduction techniques useful in implementing the invention is defined as follows. Consider a set of T training supervectors obtained from speaker-dependent models for speech recognition. Let each of these supervectors have dimension V ; thus, we can denote every supervector as $X = [x_1, x_2, \dots, x_V]^T$ (a $V \times 1$ vector). Consider a linear transformation M that can be applied to a supervector (i.e. to any vector of dimension V) to yield a new vector of dimension E (E is less than or equal to T , the number of training supervectors); each transformed vector can be denoted $W = [w_1, w_2, \dots, w_E]^T$. The values of the parameters of M are calculated in some way from the set of T training supervectors.

[0033] Thus, we have the linear transformation $W = M \cdot X$. M has dimension $E \times V$, and W has dimension $E \times 1$, where $E \leq T$; for a particular set of T training supervectors, M will be constant. Several dimensionality reduction techniques may be used to calculate a linear transformation M from a set of T training supervectors such that W has dimension $E \leq T$.

[0034] Examples include Principal Component Analysis, Independent Component Analysis, Linear Discriminant Analysis, Factor Analysis, and Singular Value Decomposition. The invention may be implemented with any such method (not only those listed) for finding such a constant linear transformation M in the special case where the input vectors are training supervectors derived from speaker-dependent modeling, and where M is used to carry out the aforementioned technique.

[0035] The basis vectors generated at step 32 define an eigenspace spanned by the eigenvectors. Dimensionality reduction yields one eigenvector for each one of the training speakers. Thus if there are T training speakers then the dimensionality reduction step 32 produces T eigenvectors. These eigenvectors define what we call eigenvoice space or eigenspace.

[0036] The eigenvectors that make up the eigenvoice space, illustrated at 34, each represent a different dimension across which different speakers may be differentiated. Each supervector in the original training set can be represented as a linear combination of these eigenvectors. The eigenvectors are ordered by their importance in modeling the data: the first eigenvector is more important than the second, which is more important than the third, and so on. Our experiments with this technique thus far show that the first eigenvector appears to correspond to a male-female dimension.

[0037] Although a maximum of T eigenvectors is produced at step 32, in practice, it is possible to discard several of these eigenvectors, keeping only the first N eigenvectors. Thus at step 36 we optionally extract N of the T eigenvectors to comprise a reduced parameter eigenspace at 38. The higher order eigenvectors can be discarded because they typically contain less important information with which to discriminate among speakers. Reducing the eigenvoice space to fewer than the total number of training speakers provides an inherent data compression that can be helpful when constructing practical systems with limited memory and processor resources.

[0038] After generating the eigenvectors from the training data each speaker in the training data is represented in eigenspace. In the case of speaker identification, each known client speaker is represented in eigenspace as depicted at step 40a and illustrated diagrammatically at 42a. In the case of speaker verification, the client speaker and potential impostor speakers are represented in eigenspace as indicated at step 40b and as illustrated at 42b. The speakers may be represented in eigenspace either as points in eigenspace (as illustrated diagrammatically in Figure 2 at 42a) or as probability distributions in eigenspace (as illustrated diagrammatically in Figure 3 at 42b).

Using the Speaker Identification or Speaker Verification System

[0039] The user seeking speaker identification or verification supplies new speech data at 44 and these data are used to train a speaker dependent model as indicated at step 46. The model 48 is then used at step 50 to construct a supervector 52. Note that the new speech data may not necessarily include an example of each sound unit. For instance, the new speech utterance may be too short to contain examples of all sound units. The system will handle this, as will

be more fully explained below.

[0040] Dimensionality reduction is performed at step 54 upon the supervector 52, resulting in a new data point that can be represented in eigenspace as indicated at step 56 and illustrated at 58. In the illustration at 58 the previously acquired points in eigenspace (based on training speakers) are represented as dots, whereas the new speech data point is represented by a star.

[0041] Having placed the new data point in eigenspace, it may now be assessed with respect to its proximity to the other prior data points or data distributions corresponding to the training speakers. Figure 4 illustrates an exemplary embodiment of both speaker identification and speaker verification.

[0042] For speaker identification, the new speech data is assigned to the closest training speaker in eigenspace, step 62 diagrammatically illustrated at 64. The system will thus identify the new speech as being that of the prior training speaker whose data point or data distribution lies closest to the new speech in eigenspace.

[0043] For speaker verification, the system tests the new data point at step 66 to determine whether it is within a predetermined threshold proximity to the client speaker in eigenspace. As a safeguard the system may, at step 68, reject the new speaker data if it lies closer in eigenspace to an impostor than to the client speaker. This is diagrammatically illustrated at 69, where the proximity to the client speaker and proximity to the closest impostor have been depicted.

The Maximum Likelihood Eigenspace Decomposition (MLED) Technique

[0044] One simple technique for placing the new speaker within eigenspace is to use a simple projection operation. A projection operation finds the point within eigenspace that is as close as possible to the point outside of eigenspace corresponding to the new speaker's input speech. It bears noting that these points are actually supervectors from which a set of HMMs can be reconstituted.

[0045] The projection operation is a comparatively crude technique that does not guarantee that the point within eigenspace is optimal for the new speaker. Furthermore, the projection operation requires that the supervector for the new speaker contain a complete set of data to represent the entire set of HMMs for that speaker. This requirement gives rise to a significant practical limitation. When using projection to constrain a new speaker to the eigenspace, that speaker must supply enough input speech so that all speech units are represented in the data. For example, if the Hidden Markov Models are designed to represent all phonemes in the English language, then the training speaker must supply examples of all phonemes before the simple projection technique can be used. In many applications this constraint is simply not practical.

[0046] The maximum likelihood technique of the invention addresses both of the above-mentioned drawbacks of simple projection. The maximum likelihood technique of the invention finds a point within eigenspace that represents the supervector corresponding to a set of Hidden Markov Models that have the maximum probability of generating the speech supplied by the new speaker.

[0047] Whereas the simple projection operation treats all members of the supervector as having equal importance, the maximum likelihood technique is based on probabilities arising from the actual adaptation data and thus tends to weight the more probable data more heavily. Unlike the simple projection technique, the maximum likelihood technique will work even if the new speaker has not supplied a full set of training data (i.e., data for some of the sound units are missing). In effect, the maximum likelihood technique takes into account the context under which the supervectors are constructed, namely from Hidden Markov Models involving probabilities that certain models are more likely than others to generate the input speech supplied by the new speaker.

[0048] In practical effect, the maximum likelihood technique will select the supervector within eigenspace that is the most consistent with the new speaker's input speech, regardless of how much input speech is actually available. To illustrate, assume that the new speaker is a young female native of Alabama. Upon receipt of a few uttered syllables from this speaker, the maximum likelihood technique will select a point within eigenspace that represents all phonemes (even those not yet represented in the input speech) consistent with this speaker's native Alabama female accent.

[0049] Figure 5 shows how the maximum likelihood technique works. The input speech from the new speaker is used to construct supervector 70. As explained above, the supervector comprises a concatenated list of speech parameters, corresponding to cepstral coefficients or the like. In the illustrated embodiment these parameters are floating point numbers representing the Gaussian means extracted from the set of Hidden Markov Models corresponding to the new speaker. Other HMM parameters may also be used. In the illustration these HMM means are shown as dots, as at 72. When fully populated with data, supervector 70 would contain floating point numbers for each of the HMM means, corresponding to each of the sound units represented by the HMM models. For illustration purposes it is assumed here that the parameters for phoneme "ah" are present but parameters for phoneme "iy" are missing.

[0050] The eigenspace 38 is represented by a set of eigenvectors 74, 76 and 78. The supervector 70 corresponding to the observation data from the new speaker may be represented in eigenspace by multiplying each of the eigenvectors by a corresponding eigenvalue, designated $W_1, W_2 \dots W_n$. These eigenvalues are initially unknown. The maximum

likelihood technique finds values for these unknown eigenvalues. As will be more fully explained, these values are selected by seeking the optimal solution that will best represent the new speaker within eigenspace.

[0051] After multiplying the eigenvalues with the corresponding eigenvectors of eigenspace 38 and summing the resultant products, an adapted model 80 is produced. Whereas the supervector of the input speech (supervector 70) may have had some missing parameter values (the "iy" parameters, for example), the supervector 80 representing the adapted model is fully populated with values. That is one benefit of the invention. Moreover, the values in supervector 80 represent the optimal solution, namely that which has the maximum likelihood of representing the new speaker in eigenspace.

[0052] The individual eigenvalues $W_1, W_2 \dots W_n$ may be viewed as comprising a maximum likelihood vector, herein referred to as maximum likelihood vector. Figure 5 illustrates vector diagrammatically at 82. As the illustration shows, maximum likelihood vector 82 comprises the set of eigenvalues $W_1, W_2 \dots W_n$.

[0053] The procedure for performing adaptation using the maximum likelihood technique is shown in Figure 6. Speech from a new speaker, comprising the observation data, is used to construct a set of HMMs as depicted at 100. The set of HMMs 102 is then used in constructing a supervector as depicted at 104. As illustrated, the supervector 106 comprises a concatenated list of HMM parameters extracted from the HMM models 102.

[0054] Using the supervector 106, a probability function Q is constructed at 108. The presently preferred embodiment employs a probability function that represents the probability of generating the observed data for the pre-defined set of HMM models 102. Subsequent manipulation of the probability function Q is made easier if the function includes not only a probability term P but also the logarithm of that term, log P.

[0055] The probability function is then maximized at step 110 by taking the derivative of the probability function individually with respect to each of the eigenvalues $W_1, W_2 \dots W_n$. For example, if the eigenspace is of dimension 100, this system calculates 100 derivatives of the probability function Q setting each to zero and solving for the respective W. While this may seem like a large computation, it is far less computationally expensive than performing the thousands of computations typically required of conventional MAP or MLLR techniques.

[0056] The resulting set of Ws, so obtained, represent the eigenvalues needed to identify the point in eigenspace corresponding to the point of maximum likelihood. Thus the set of Ws comprises a maximum likelihood vector in eigenspace. In this regard, each of the eigenvectors (eigenvectors 74, 76 and 78 in Fig. 5) define a set of orthogonal vectors or coordinates against which the eigenvalues are multiplied to define a point constrained within eigenspace. This maximum likelihood vector, depicted at 112, is used to construct supervector 114 corresponding to the optimal point in eigenspace (point 66 in Fig. 4). Supervector 114 may then be used at step 116 to construct the adapted model 118 for the new speaker.

[0057] In the context of the maximum likelihood framework of the invention, we wish to maximize the likelihood of an observation $O=o_1 \dots o_T$ with regard to the model λ . This may be done by iteratively maximizing the auxiliary function Q (below); where λ is the current model at the iteration and $\hat{\lambda}$ is the estimated model. We have:

$$Q(\lambda, \hat{\lambda}) = \sum_{\theta \in \text{states}} P(O, \theta | \lambda) \log [P(O, \theta | \hat{\lambda})]$$

[0058] As a preliminary approximation, we might want to carry out a maximization with regards to the means only. In the context where the probability P is given by a set of HMMs, we obtain the following:

$$Q(\lambda, \hat{\lambda}) = \text{const} - \frac{1}{2} P(O | \lambda) \sum_{\substack{\text{states} \\ \text{in } \lambda}} \sum_{\substack{\text{mix} \\ \text{in } \lambda}} \sum_{\substack{\text{time} \\ t}} \{ \gamma_m^{(s)}(t) [n \log(2\pi) + \log |C_m^{(s)}| + h(o_t, m, s)] \}$$

where:

$$h(o_t, m, s) = (o_t - \hat{\mu}_m^{(s)})^T C_m^{(s)-1} (o_t - \hat{\mu}_m^{(s)})$$

and let:

O_t be the feature vector at time t

$C_m^{(s)-1}$ be the inverse covariance for mixture gaussian m of state s
 $\hat{\mu}_m^{(s)}$ be the approximated adapted mean for state s, mixture component m
 $\gamma_m^{(s)}(t)$ be the P(using mix gaussian m| λ, o_t)

5 [0059] Suppose the gaussian means for the HMMs of the new speaker are located in eigenspace. Let this space be spanned by the mean supervectors $\bar{\mu}_j$ with $j=1 \dots E$,

$$\bar{\mu}_j = \begin{bmatrix} \bar{\mu}_1^{(1)}(j) \\ \bar{\mu}_2^{(1)}(j) \\ \vdots \\ \bar{\mu}_m^{(s)}(j) \\ \vdots \\ \bar{\mu}_{M_{S_1}}^{(S_1)}(j) \end{bmatrix}$$

where $\bar{\mu}_m^{(s)}(j)$ represents the mean vector for the mixture gaussian m in the state s of the eigenvector (eigenmodel) j.

[0060] Then we need:

$$\hat{\mu} = \sum_{j=1}^E w_j \bar{\mu}_j$$

30 [0061] The $\bar{\mu}_j$ are orthogonal and the w_j are the eigenvalues of our speaker model. We assume here that any new speaker can be modeled as a linear combination of our database of observed speakers. Then

$$\hat{\mu}_m^{(s)} = \sum_{j=1}^E w_j \bar{\mu}_m^{(s)}(j)$$

with s in states of λ , m in mixture gaussians of M.

[0062] Since we need to maximize Q, we just need to set

$$\frac{\partial Q}{\partial w_e} = 0, \quad e=1 \dots E.$$

45 (Note that because the eigenvectors are orthogonal, $\frac{\delta w_l}{\delta w_j} = 0, l \neq j$.)
Hence we have

$$\frac{\partial Q}{\partial w_e} = 0 = \sum_{\substack{\text{states} \\ \text{in } \lambda}} \sum_{\substack{\text{mix} \\ \text{gauss} \\ \text{in } S}} \sum_{t=1}^T \left\{ \frac{\partial}{\partial w_e} \gamma_m^{(s)}(t) h(o_t, s) \right\}, \quad e=1 \dots E.$$

Computing the above derivative, we have:

$$0 = \sum_s \sum_m \sum_l \gamma_m^{(s)}(l) \left\{ -\bar{\mu}_m^{(s)T}(e) C_m^{(s)-1} o_l + \sum_{j=1}^E w_j \bar{\mu}_m^{(s)T}(j) C_m^{(s)-1} \bar{\mu}_m^{(s)}(e) \right\}$$

from which we find the set of linear equations

$$\sum_s \sum_m \sum_l \gamma_m^{(s)}(l) \bar{\mu}_m^{(s)T}(e) C_m^{(s)-1} o_l = \sum_s \sum_m \sum_l \gamma_m^{(s)}(l) \sum_{j=1}^E w_j \bar{\mu}_m^{(s)T}(j) C_m^{(s)-1} \bar{\mu}_m^{(s)}(e), \quad e=1 \dots E.$$

Assessing Proximity in Eigenspace

[0063] When representing speakers as points in eigenspace, a simple geometric distance calculation can be used to identify which training data speaker is closest to the new speaker. When representing speakers as distributions in eigenspace, proximity is assessed by treating the new speaker data as an observation O and by then testing each distribution candidate (representing the training speakers) to determine what is the probability that the candidate generated the observation data. The candidate with the highest probability is assessed as having the closest proximity. In some high-security applications it may be desirable to reject verification if the most probable candidate has a probability score below a predetermined threshold. A cost function may be used to thus rule out candidates that lack a high degree of certainty.

[0064] Assessing the proximity of the new speaker to the training speakers may be carried out entirely within eigenspace, as described above. Alternatively, a Bayesian estimation technique can be used for even greater accuracy.

[0065] To enhance the proximity assessment using Bayesian estimation, the Gaussian densities of the training speakers within eigenspace are multiplied by the estimated marginal density in the orthogonal complement space that represents the speaker data that were discarded through dimensionality reduction. In this regard, recognize that performing dimensionality reduction upon the speaker model supervectors results in a significant data compression from high-dimensionality space to low-dimensionality space. Although dimensionality reduction preserves the most important basis vectors, some higher-order information is discarded. The Bayesian estimation technique estimates a marginal Gaussian density that corresponds to this discarded information.

[0066] To illustrate, assume that the original eigenspace is constructed by linear transformation of the supervector through a dimensionality reduction process whereby M components are extracted from the larger number N of all components. The smaller extracted M components represent a lower-dimensional subspace of the transformation basis that correspond to the maximal eigenvalues. Thus the eigenspace is defined by components $i=1 \dots M$, whereas the discarded minor components correspond to $i=M+1 \dots N$. These two sets of components define two mutually exclusive and complementary subspaces, the principal subspace represents the eigenspace of interest and its orthogonal component represents the data that were discarded through dimensionality reduction.

[0067] We can compute the likelihood estimate as the product of the Gaussian densities in these two respective orthogonal spaces by the following equation:

$$\hat{P}(x|\Omega) = P_E(x|\Omega) \cdot P_{\bar{E}}(x|\Omega)$$

In the above equation, the first term is the single Gaussian density in eigenspace E and the second term is the single Gaussian distribution in the space orthogonal to the eigenspace. It turns out that both terms can be estimated entirely from the set of training data vectors, using only the projections into eigenspace and the residuals.

Further Embodiments

[0068] In the foregoing examples, speech has been represented as Hidden Markov Models (HMMs). Hidden Markov Models are employed in popular use today in many speech recognizers and thus they may be used for speaker verification and speaker identification purposes as well. However, the techniques of the present invention are not limited to the use of Hidden Markov Models. For example, a useful and effective system for speaker verification and/or speaker identification may be implemented using Gaussian Mixture Models (GMMs). Gaussian Mixture Models are single state models that may be trained on text-independent or text-dependent training data. By comparison, typically Hidden Markov Models have plural states and are trained on speech data that have been labeled according to the text used for the training data. Gaussian Mixture Models may therefore be seen as a special case of Hidden Markov Models, in which

only a single state is used and in which the training data need not be labeled.

[0069] The Gaussian Mixture Model (GMM) may be used for speaker identification and verification purposes by assigning individual component Gaussians to represent broad acoustic classes. These classes may represent general speaker-dependent vocal tract configurations that are useful for modeling speaker identity. The Gaussian mixture density provides a smooth approximation to the underlying long term sample distribution of observations obtained from utterances by a given speaker. See Reynolds, D. A., "Speaker Identification and Verification Using Gaussian Mixture Speaker Models", Speech Communication, Vol. 17, pp. 91-108, 1995.

[0070] A Gaussian mixture density is a weighted sum of M component densities and is given by the equation,

$$f(\bar{x}) = \sum_{i=1}^M p_i b_i(\bar{x})$$

where \bar{x} is a D — dimensional vector,
 $i = 1, \dots, M$ are the component densities and
 $p_i, i = 1, \dots, M$, are the mixture weights,
 Each component density is a D-variate Gaussian function of the form,

$$b_i(\bar{x}) = \frac{1}{(2\pi)^{D/2} |\Sigma_i|^{1/2}} e^{-\frac{1}{2} (\bar{x} - \mu_i)^T \Sigma_i^{-1} (\bar{x} - \mu_i)}$$

with the mean vector μ_i and covariance matrix Σ_i . The mixture weights further satisfy the constraint that

$$\sum_{i=1}^M p_i = 1$$

The complete GM density is parameterized by the mean vector, covariance matrices and mixture weights from all the component densities.

$$\lambda = \{p_i, \mu_i, \Sigma_i\}, i=1, \dots, M$$

[0071] Moreover, it will be appreciated that while HMMs and GMMS have been illustrated here, other types of speech models may also be used. The best models for this purpose are those that are numerically expressed (e.g., as floating point numbers) such that a speaker space can be defined mathematically. For illustration purposes, a GMM Model 120 has been illustrated in Figure 7.

[0072] In the foregoing examples, speaker space has been represented as a linear combination of eigenvoices. However, the techniques of the present invention are not limited to a speaker space of this type. More generally, speaker space is a set of mathematical constraints derived from a set of training speakers and representing *a priori* knowledge which must be satisfied by new speakers. Besides eigenvoice based speaker space, other approaches include (but are not limited to) "reference speaker weighting" (See Hazen, T.J., and Glass, J.R., "A Comparison of Novel Techniques for Instantaneous Speaker Adaptation", Eurospeech Proceedings, pp 2047-50, 1997) and speaker clustering (See Kosaka, T., and Sagayama, S., "Tree-Structured Speaker Clustering for Fast Speaker Adaptation", ICASSP pp. 1-245 to 1-248, 1994).

[0073] Figure 8 illustrates that there are also alternatives to constructing the speaker space, representing the enrollment speech in the speaker space, and determining whether the test speaker is one of the client speakers. At the outset, an important consideration regarding construction of the speaker space relates to selection of the training speakers 122. While client speakers 124 could be used to collect the training data 22, certain benefits can be achieved by using a second group of individuals as the training speakers 122. For example, this approach allows the training speakers 122 to be arbitrarily large and will generally allow more diverse training data. For example, paid individuals or volunteers could be selected in advance from a population much larger than the client speaker population. The

selected population of training speakers could have no specific relationship to the population defined by the client speakers 124 (other than their ability to speak). These training speakers would each provide a comparatively large sample of training speech. This would allow the construction of a much more diverse speaker space, based on comparatively welltrained speech models, and would allow a significant reduction in the amount of data required from the client speakers 124. Thus, at the client enrollment step, only a few seconds of speech from each client would be required as opposed to several minutes of speech. This is a key advantage of the speaker space approach.

[0074] Step 132 illustrates the process of training the speaker space. As discussed above, the result can be a set of either (preferably text-independent) GMM speech models as illustrated at 126 or text-dependent speech models as discussed above. Thus, although Gaussian Mixture Models have been illustrated here, the invention is not restricted to Gaussian Mixture Models (or Hidden Markov Models, for that matter). Rather, any speech model having parameters suitable for concatenation may be used.

[0075] The speech models may optionally be fine-tuned or adapted to account for differences between the environment used during training and the environment that will be used during subsequent use for speaker verification and/or speaker identification. Typically, the training data are collected under controlled conditions (known background noise qualities, standardized microphones and signal processing equipment, controlled microphone placement, etc.). In use, the system may be deployed in an office environment, for example, where the environmental conditions are quite different from those of the training environment. To accommodate such variation, an environmental adaptation process may be used to improve the training speaker models for their specific use in a given environment. MLLR adaptation may be used for this purpose. Other known adaptation techniques may be used as well.

[0076] In the currently preferred embodiment, models for each speaker are used to create a supervector. The supervector may be formed by concatenating the parameters of the model for each speaker. Where Gaussian Mixture Models are used, the floating point numbers used to represent the Gaussian mixtures may be concatenated for each speaker.

[0077] After constructing the supervectors, a technique that reduces the number of degrees of freedom in a speech model for a particular speaker is applied. Such techniques operate on training speaker data to generate a speaker space of reduced dimensionality. While any such technique can be used, linear discriminant analysis (LDA) is shown here and is presently preferred. Thus, in addition to the supervectors, step 132 makes use of global within-speaker scatter matrix data 130. This is important to note because this type of data is generally not part of a speaker's speaker-dependent model.

[0078] After a technique such as PCA or LDA has generated an initial set of basis vectors 134, an optional step 136 of reestimating the speaker space can be performed. Here, a technique such as MLES may rotate the basis vectors 134 in space such that the likelihood of the training data according to the training speaker models in the space is maximized. The result would be an improved set of basis vectors 138. Details of the MLES technique are provided below.

[0079] After the speaker space has been generated, the system may be used to enroll one or more client speakers so that speaker identification and/or speaker verification may be performed with respect to those client speakers. Enrollment is performed at step 140 where each client speaker is represented in speaker space based on a short utterance of enrollment speech. This is done by training an enrollment speech model upon the enrollment speech from the client speaker (possibly as little as a few words), and then placing the client speakers into the speaker space by MLED or projection, as discussed above. If desired, speaker or environment adaptation techniques such as MLLR may be employed to improve the speech models of one or more client speakers, or to re-estimate the speaker space so that it better models the new environment (i.e., the environment in which the client speakers were recorded).

[0080] At this point, the power of the speaker space may be more fully appreciated. When the client speaker provides a very short sample of speech, there may not be enough data to construct a full speech model for that speaker. Nevertheless, by placing the partial model into the speaker space at its proper location as dictated by the MLED procedure (or through projection), the speaker space will fill in the details, allowing a complete speech model for that speaker to be later generated.

[0081] After the speaker space has been generated and all client speakers enrolled, the system is ready for use. To perform speaker verification or speaker identification on a test speaker, a sample of speech is taken from that speaker and assessed using the client-populated speaker space. In the preceding examples, speaker verification and speaker identification were performed by placing the speech of the test speaker into the speaker space, to determine, through suitable distance measure, which client speaker the test speaker was closest to. The following will describe an alternate technique.

[0082] Instead of placing the test speaker into the speaker space, the alternate technique expands the client speaker vector points within the speaker space back into complete speech models. Recall that even though the initial client enrollment speech was very short (quite possibly resulting in incomplete speech models) the points in speaker space will generate complete speech models. This is so because the original speaker space contains a great deal of *a priori* knowledge about the characteristics of human speech. In other words, only a few spoken words from a client speaker

are sufficient to place that client speaker into the speaker space, where a full and complete speech model may be inferred.

[0083] In the alternate technique, each client speaker point within the speaker space is used to generate its corresponding complete speech model. Then, each of the client speaker models is assessed vis-à-vis speech from the test speaker. The client model with the highest probability of producing the test speech is then used for speaker identification and/or speaker verification purposes.

[0084] Expansion of the client speaker's position in speaker space vectors back into speech models is shown at step 144 in Figure 8. Specifically, the corresponding complete speech models 146 are generated from their locations in the speaker space 142. These models are then used for subsequent speaker verification and/or speaker identification. Each of the models is tested vis-à-vis the test speech data supplied by a test speaker (user of the system). The model that has the highest probability of generating the test speech is used for subsequent speaker verification and identification purposes. Figure 9 diagrammatically illustrates the process where speaker models 146 are used to assess test speech. At step 148 speech data supplied by test speaker(s) 150 are submitted to the probabilistic client speaker models 146 as part of a likelihood analysis. Each test speaker is assigned to whichever client yields the highest likelihood of producing his or her speech; alternatively, the test speaker may be classified as an impostor. Thus, the final assessment does not occur in speaker space but in model space.

[0085] It is also worth noting that the speaker space can be adapted as new speech is obtained during client enrollment. In the event the client environment differs from the original training environment (as often it will) environmental adaptation can be performed. For example, because the training-derived speaker space produces models that represent or express inter-speaker variability, these models can be used to estimate an environment mismatch function and to apply that function to the speaker space (i.e. as a linear transformation). This would prevent irrelevant features of the test environment from interfering with speaker verification and identification.

Enrollment of Difficult Client Speakers

[0086] Even though the present invention generally requires very little enrollment speech, it may be beneficial to overall performance to ask a minority of troublesome (i.e. variable) clients for more enrollment data. This is done because, in practice, it is typically a small set of specific client speakers that cause most of the misidentifications. In the present approach, these clients are identified while they are enrolling and more speech is required of those clients. In other words, the enrollment speech model can be readily trained upon additional speech from the client speaker when the enrollment speech meets predetermined conditions. For example, Figure 10 shows that the predetermined conditions can be defined to include the location of a first utterance in the speaker space being a predetermined distance from the location of a second utterance in the speaker space. If the average within-speaker distance is much greater than the average of the two locations, the client in question is therefore asked for more enrollment data. As shown in Figure 11, the predetermined conditions could also be defined to include the first location being positioned in an area of the speaker space having a predetermined density (i.e. a "crowded" area). In this case, an approach such as MLED allows prior information about the distribution of speakers in the speaker space to be taken into account. This distribution could be estimated from the training data or from the enrollment data.

MLES Speaker Space Re-estimation

[0087] As presented above, the speaker space technique confines the speaker models to a very low-dimensional linear vector space, called the speaker space. This speaker space summarizes *a priori* knowledge about the speaker models obtained during initial system training. While the speaker space in its initially-generated form will serve as a powerful tool for speaker identification and speaker verification, as discussed above, additional improvements to the speaker space are possible through a technique called Maximum-Likelihood Eigenspace (MLES). The MLES method performs re-estimation on the training data. It results in the vectors within the speaker space being rotated, such that the likelihood of the training data according to the training speaker models in the space is maximized. The MLES technique begins by integrating values as hidden data in the estimation problem, yielding:

$$\hat{M} = \arg \max_M \sum_{q=1}^T \int \log L(O, w | M) P_o(w, q) dw$$

where $P_o(w, q)$ contains prior information about speaker q (e.g., the probability of a person of a given dialect or sex to appear). It is extensively used for unbalanced sets of speakers. For instance, we may set for a given k

$$P_o(w_k, q) = \begin{cases} 1 & \text{if } w_k > 0 \quad \text{and } q^{\text{th}} \text{ speaker is male} \\ 1 & \text{if } w_k < 0 \quad \text{and } q^{\text{th}} \text{ speaker is female} \\ 0 & \text{elsewhere} \end{cases}$$

Seed speaker voices can be obtained through PCA, linear discriminant analysis (LDA), speaker clustering, or can be given as a set of speaker dependent models. When no particular knowledge about w_k is known, we use MLED to replace the integration operator by a maximum operator.

[0088] The re-estimation formula is relatively easy to derive

$$\bar{\mu}_e^{(m)} = \frac{\sum_q L_q w_q^{(e)} \sum_t \gamma_m(t) \{p_t - \bar{\mu}_q^{(m)}(e)\}}{\sum_q L_q (w_q^{(e)})^2 \sum_t \gamma_m(t)}$$

where q, m, e represent a speaker, a distribution, and a speaker space basis vector. L_q is the posterior probability of the utterances $O^{(q)}$ of the speaker, $L_q \gamma_m(t)$ is the observed posterior probability. $w_q^{(e)}$ is the current estimate of the e^{th} coordinate of speaker q . Finally, $\bar{\mu}_q^{(m)}$ is the complement of the estimated mean, i.e.,

$$\bar{\mu}_q^{(m)}(e) = \sum_{k=1, k \neq e}^E w_q^{(k)} \bar{\mu}_k^{(m)}, e = 1, \dots, E$$

[0089] From the foregoing it will be appreciated that the invention provides powerful techniques for performing speaker verification and/or speaker identification. While several examples of the invention have been illustrated here, it will be appreciated by those of skill in this art that numerous other variations are possible within the scope of the appended claims.

Claims

1. A method for assessing speech with respect to a predetermined client speaker, the method comprising the steps of:

training a set of speech models upon the speech from a plurality of training speakers;
constructing a speaker space from the set of speech models to represent said plurality of training speakers;
representing enrollment speech from said client speaker as a first location in said speaker space; and
determining whether a new speaker is the client speaker based on the first location and new speech data from the new speaker.

2. The method of claim 1 further including the steps of:

generating a probabilistic speech model from the first location; and
assessing a likelihood between the probabilistic speech model and the new speech data and using said assessment as an indication of whether the new speaker is the client speaker.

3. The method of claim 1 further including the steps of:

using said speech data to generate a representation of the new speaker as a second location in the speaker space; and
assessing the proximity between the first and second locations and using the assessment as an indication of whether the new speaker is the client speaker.

4. The method of claim 1 further including the steps of:

training an enrollment speech model upon the enrollment speech from the client speaker; and
generating a representation of the client speaker as the first location in the speaker space.

5. The method of claim 4 further including the step of training the enrollment speech model upon additional speech from the client speaker when the enrollment speech meets predetermined conditions.
6. The method of claim 5 wherein the enrollment speech includes a first utterance and a second utterance, the method further including the step of defining the predetermined conditions to include the location of the first utterance in the speaker space being a predetermined distance from the location of the second utterance in the speaker space.
7. The method of claim 5 further including the step of defining the predetermined conditions to include the first location being positioned in an area of the speaker space having a predetermined density.
8. The method of claim 1 further including the step of training a set of text-independent speech models.
9. The method of claim 1 further including the step of training a set of text-dependent speech models.
10. The method of claim 1 further including the steps of:
 - obtaining a speaker-dependent supervector for each said training speaker; and
generating a speaker space having lower dimensionality than the speech models based on the speaker dependent supervectors.
11. The method of claim 10 further including the steps of:
 - obtaining global within-speaker scatter matrix data; and
performing dimensionality reduction based on the matrix data.
12. The method of claim 1 further including the step of reestimating the speaker space.
13. The method of claim 1 further including the step of adapting the speaker space based on information regarding a client enrollment environment.
14. The method of claim 1 further including the step of assessing the proximity between said first and second locations and using said assessment as an indication of whether the new speaker is the client speaker.
15. The method of claim 1 wherein the plurality of training speakers includes the client speaker.
16. The method of claim 1 wherein the plurality of training speakers does not include the client speaker.
17. The method of claim 1 further including the step of constructing the speaker space by reducing the number degrees of freedom for each training speaker's speech model.
18. The method of claim 1 further including the step of performing speaker identification.
19. The method of claim 1 further including the step of performing speaker verification by deciding whether the new speaker is the client speaker or an imposter.




FIG. 1




FIG. 2




FIG. 3

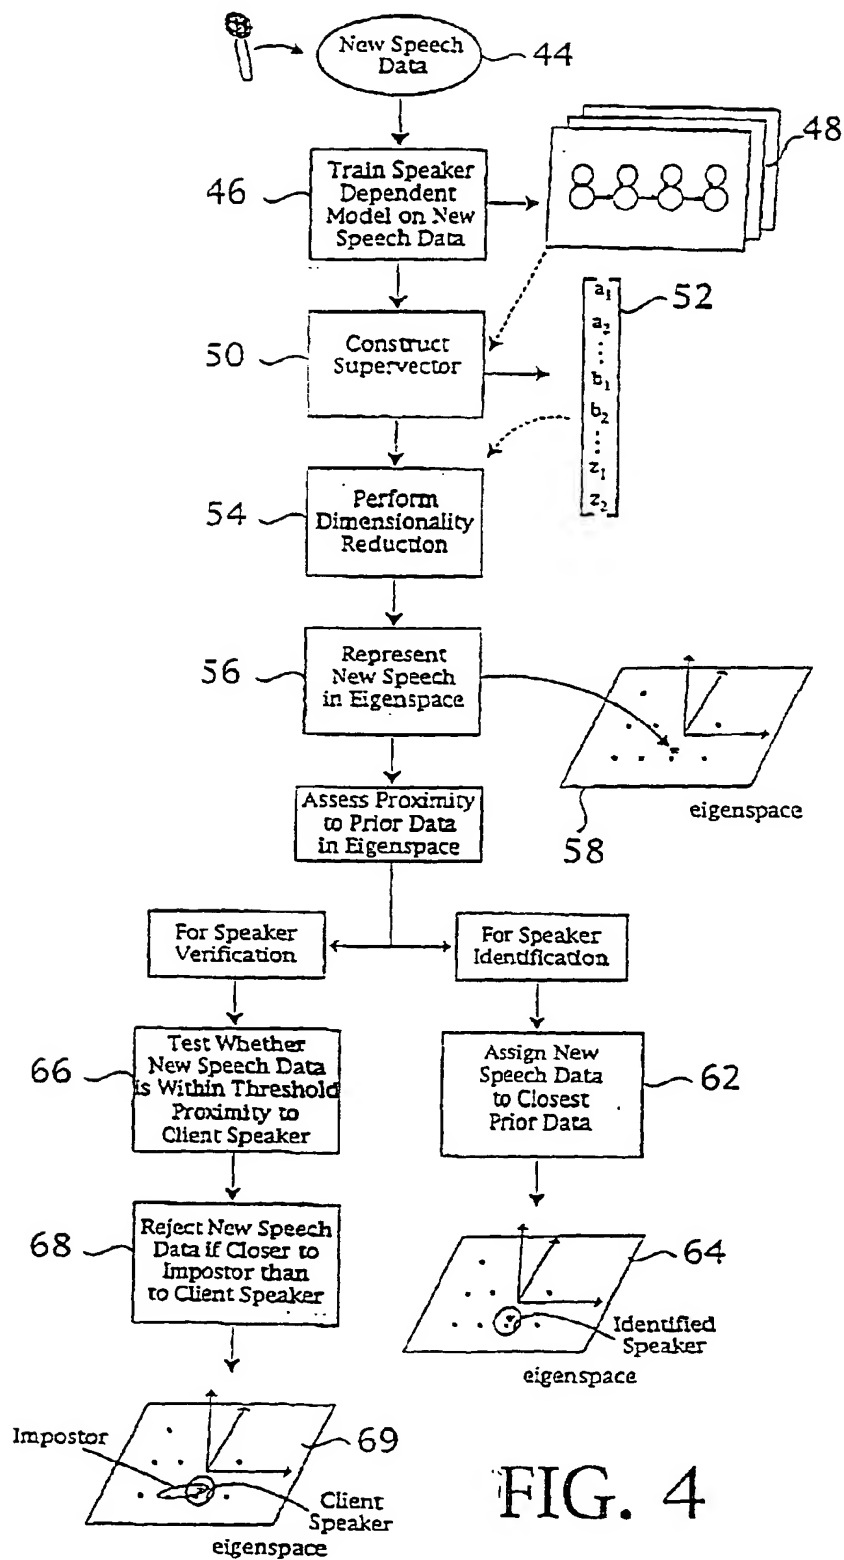


FIG. 4




FIG. 5

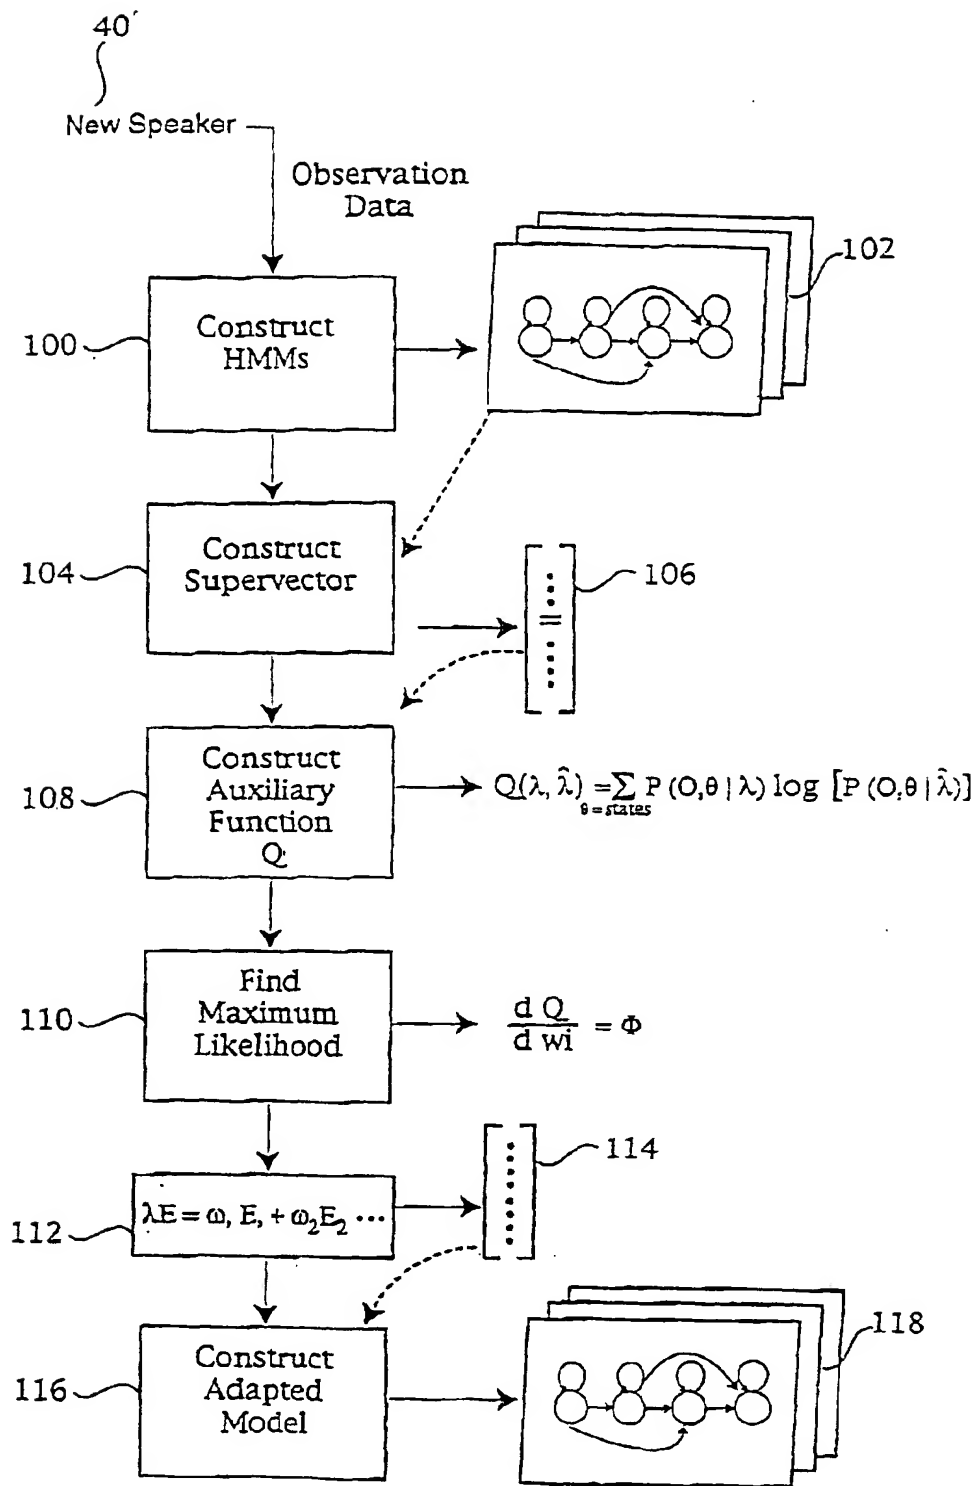


FIG. 6

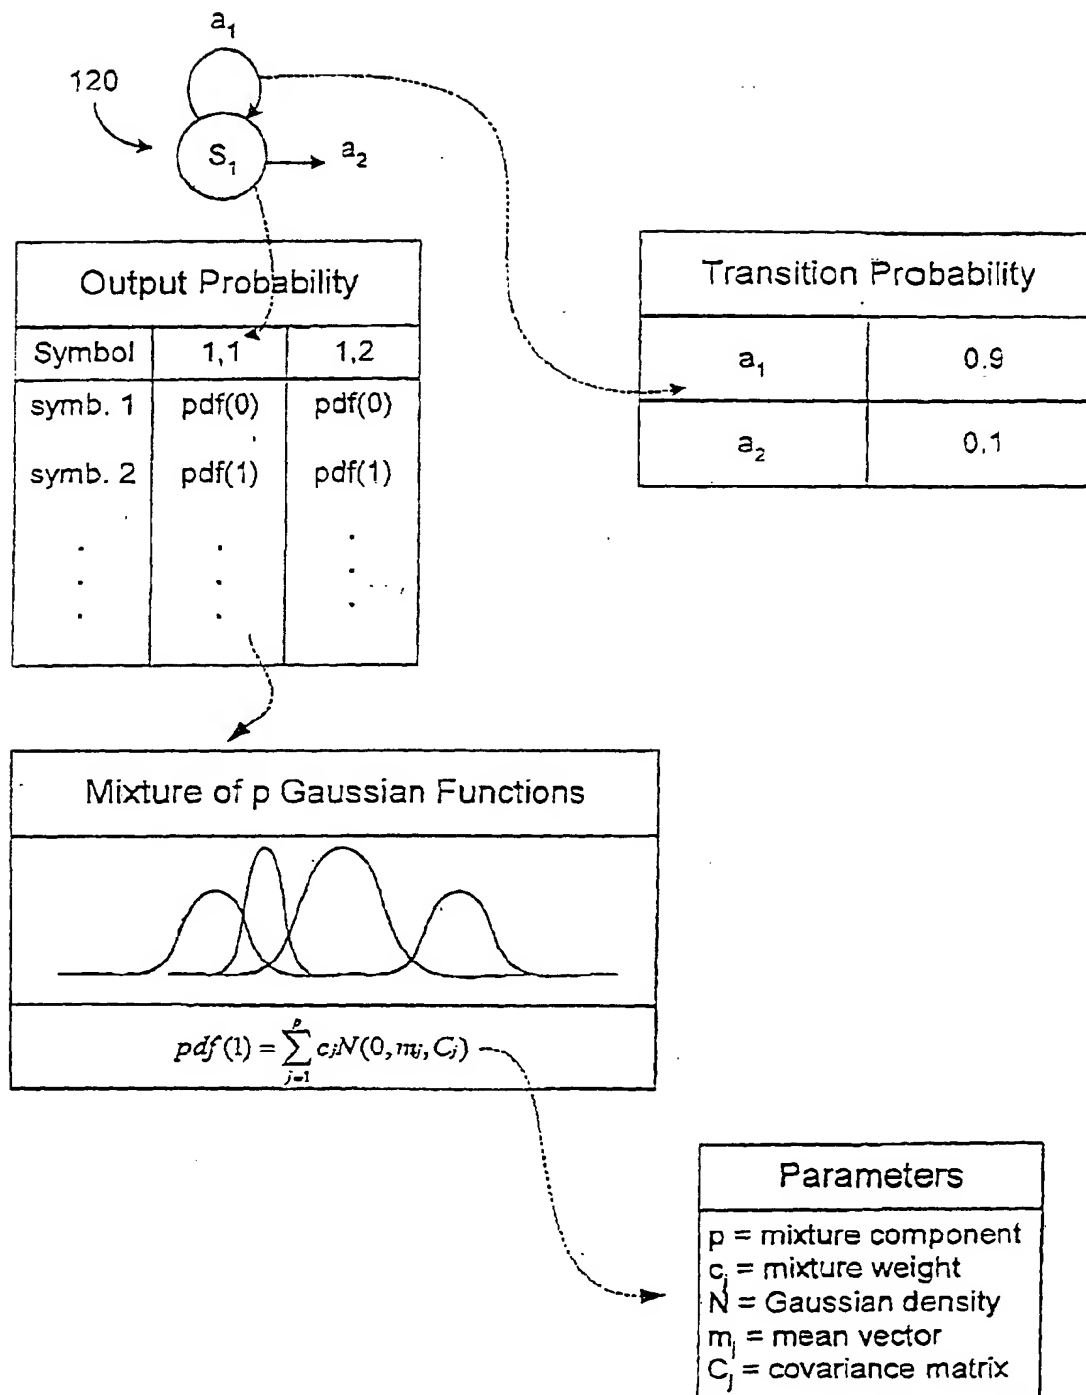


FIG. 7




FIG. 8




FIG. 9

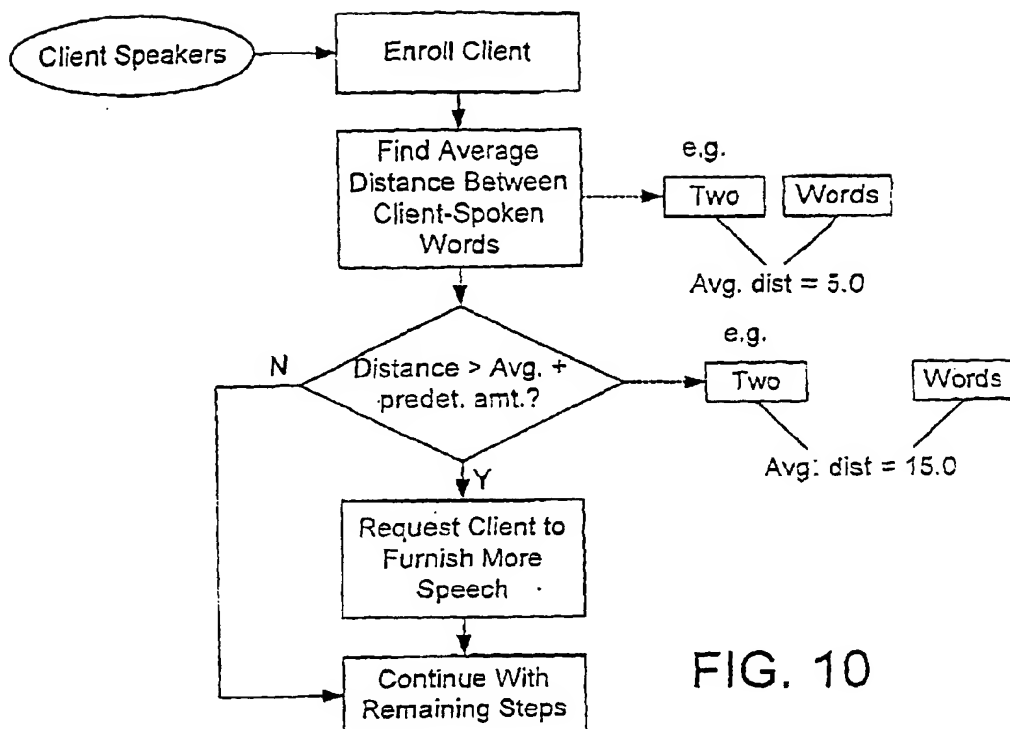


FIG. 10

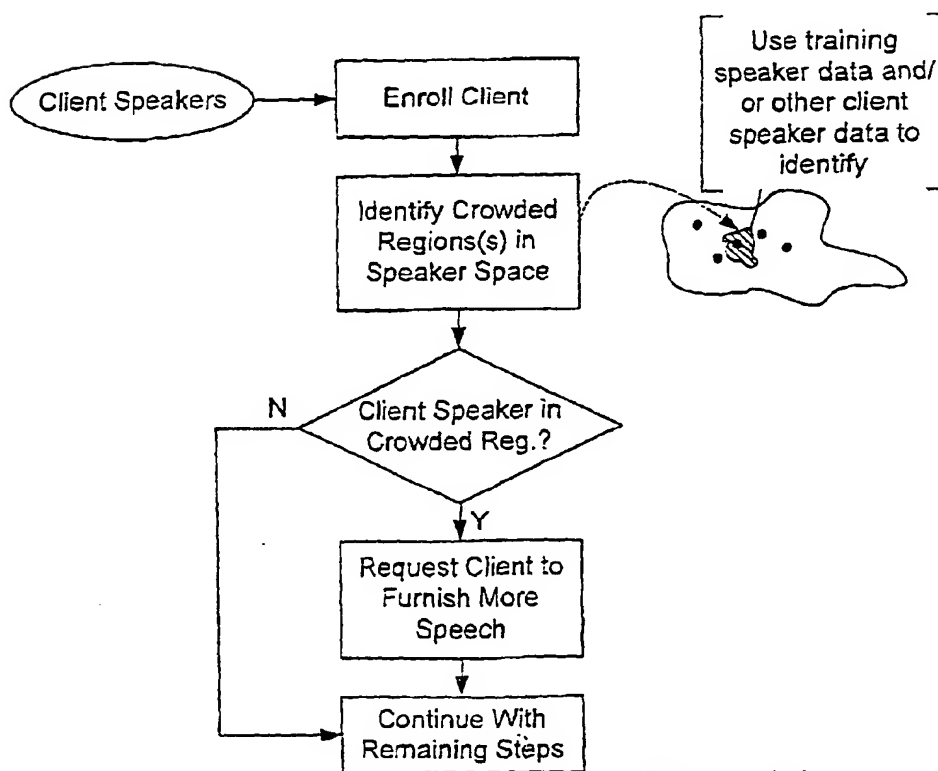


FIG. 11



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 5725

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	EP 0 984 431 A (MATSUSHITA ELECTRIC IND CO LTD) 8 March 2000 (2000-03-08)	1-4, 8-10, 14-19	G10L17/00
Y	* the whole document *	5-7, 11, 13	
Y	EP 0 953 968 A (MATSUSHITA ELECTRIC IND CO LTD) 3 November 1999 (1999-11-03) * abstract *	13	
Y	FREDOUILLE C ET AL: "Behavior of a Bayesian adaptation method for incremental enrollment in speaker verification" 2000 IEEE INTERNATIONAL CONFERENCE ON ACOUSTICS, SPEECH, AND SIGNAL PROCESSING. PROCEEDINGS (CAT. NO.00CH37100), ISTANBUL, TURKEY, 5-9 JUNE 2000, pages II1197-II1200 vol.2, XP002181932 2000, Piscataway, NJ, USA, IEEE, USA ISBN: 0-7803-6293-4 * abstract *	5-7	
Y	CHENGJUN LIU ET AL: "Enhanced Fisher linear discriminant models for face recognition" PROCEEDINGS. FOURTEENTH INTERNATIONAL CONFERENCE ON PATTERN RECOGNITION (CAT. NO.98EX170), BRISBANE, QLD., AUSTRALIA, 16-20 AUG. 1998, pages 1368-1372 vol.2, XP002181933 1998, Los Alamitos, CA, USA, IEEE Comput. Soc, USA ISBN: 0-8186-8512-3 * abstract *	11	
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 5 November 2001	Examiner: Quélavoine, R
CATEGORY OF CITED DOCUMENTS X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document		T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document	

LPU FORM 1503 03-92 (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 5725

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	<p>NGUYEN P ET AL: "EIGENVOICES: A COMPACT REPRESENTATION OF SPEAKERS IN MODEL SPACE VOIX PROPRES: VERS UNE REPRESENTATION COMPACTE DES LOCUTEURS DANS L'ESPACE DES MODELES"</p> <p>ANNALES DES TELECOMMUNICATIONS - ANNALS OF TELECOMMUNICATIONS, PRESSES POLYTECHNIQUES ET UNIVERSITAIRES ROMANDES, LAUSANNE, CH, vol. 55, no. 3/4, March 2000 (2000-03), pages 163-171, XP000948708 ISSN: 0003-4347</p> <p>* the whole document *</p>	1-4, 8-10, 14-19	
			TECHNICAL FIELDS SEARCHED (Int.Cl.7)
The present search report has been drawn up for all claims			
Place of search THE HAGUE		Date of completion of the search 5 November 2001	Examiner Quélavoine, R
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03/02 (P04001)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 30 5725

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report. The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

05-11-2001

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
EP 0984431 A	08-03-2000	US 6141644 A	31-10-2000
		CN 1247363 A	15-03-2000
		EP 0984431 A2	08-03-2000
		JP 2000081894 A	21-03-2000
EP 0953968 A	03-11-1999	US 6263309 B1	17-07-2001
		EP 0953968 A2	03-11-1999
		JP 11338491 A	10-12-1999

EPC FORM P/459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)